A Testbed for Characterizing Dynamic Response of Virtual Environment Spatial Sensors

Bernard D. Adelstein*, Eric R. Johnston*, and Stephen R. Ellis

Aerospace Human Factors Research Division
NASA Ames Research Center
Moffett Field, CA 94035

*Sterling Software

Keywords: virtual environments, input devices, spatial sensors, system calibration, sensor lag

ABSTRACT

This paper describes a testbed and method for characterizing the dynamic response of the type of spatial displacement transducers commonly used in virtual environment (VE) applications. The testbed consists of a motorized rotary swing arm that imparts known displacement inputs to the VE sensor. The experimental method involves a series of tests in which the sensor is displaced back and forth at a number of controlled frequencies that span the bandwidth of volitional human movement. During the tests, actual swing arm angle and reported VE sensor displacements are collected and time stamped. Because of the time stamping technique, the response time of the sensor can be measured directly, independent of latencies in data transmission from the sensor unit and any processing by the interface application running on the host computer. Analysis of these experimental results allows sensor time delay and gain characteristics to be determined as a function of input frequency. Results from tests of several different VE spatial sensors (Ascension, Logitech, and Polhemus) are presented here to demonstrate use of the testbed and method.

INTRODUCTION

If virtual environment (VE) technology is to be employed effectively and reliably in experimental human factors research, scientific data visualization, or human operator training, it must be subjected to the same rigorous performance characterization procedures as any other engineering or scientific research apparatus (Ellis, 1991), (Durlach, Pew, Aviles, DiZio, and Zeltzer, 1992). Objective quantitative understanding of VE system characteristics would permit assessment of the quality of experimental and analytic results, and prediction of the expected benefits of training obtained by VE techniques. For developers of VE hardware and software, detailed quantitative descriptions of VE systems and components will enable optimization to maximize performance. All consumers of VE hardware and software stand to profit from critical objective data on how well one component or configuration rates against another.

Adelstein, B. D., Johnston, E. R., & Ellis, S. R. (1992). A testbed for characterizing the response of virtual environment spatial sensors. *The 5th Annual ACM Symposium on User Interface Software and Technology*. Monterey, California. ACM. 15-22.

To date, characterizations have been developed for both spatial and temporal properties of a number of VE systems and components. The rational first approach to assessing VE performance has been to treat temporal and spatial properties as being separable—*i.e.*, spatial characteristics are considered time-invariant, and dynamic characteristics are assumed to produce purely temporal phenomena that have no spatial element. For example, distortions in the static spatial characteristics of commercially available VE displacement transducers have been examined by Hirose, Kijima, Sato, and Ishii (1990), Burdea, Dunn, Immendorf, and Mallik (1991), and Bryson (1992). Timing analyses of VE systems have been reported by Hirose *et al.* (1990), Bryson and Fisher (1990), and Liang, Shaw, and Green (1991).

Overall system latency—the time elapsed from motion of the user's instrumented hand or head mounted display until representation of that movement in the display—remains as one of the more widely acknowledged shortcomings of current VE technology. For the human interacting with the virtual environment, excessive system latency can trigger "simulator sickness" due to the sensory mismatch between signals in ocular and vestibular pathways (Oman, 1991). Time delay can introduce conflicts between visual input and motor output when performing manual tasks (Held, Efsathiou, and Greene, 1966). Sheridan and Ferrell (1963) demonstrated that human subjects slow their manual response time to preserve spatial acuity in the presence of system delays. It is also well known from systems and control engineering principles (*e.g.*, Ogata, 1970) that, in general, time lag reduces the stability and effective bandwidth of feedback control systems.

Overall latency in the typical VE is the sum of hardware and software processing times inherent in the components that form the complete system. These major components include the transducers that accept human input (*e.g.*, manual, voice), the simulation engine that generates the VE and governs human interaction with it, and the output rendering equipment (*e.g.*, visual, aural, kinesthetic displays). Of equal significance is the finite communication time required to transfer information between these components.

Latency is distinct from update period. Various update periods or "sample and hold" intervals can be present in VE systems. They include the time elapsed between successive samples of human input by transducers, the interval between successive transfers of processed transducer measurements to the simulation engine, the cycle time for the simulation engine to recalculate elements in the VE, and the rate at which the output display is re-rendered. Since sampling in the various system components usually does not occur synchronously, the effective update period of a complete VE system can be considered to be the time for the slowest component to be refreshed. While

- The testbed must be able to apply displacement inputs over a range of frequencies. Since the intended use of VE spatial sensors is to monitor human movement, test frequencies should span the bandwidth of *volitional* human limb and head motion—typically up to 3 Hz for all but the smallest amplitude hand movements (Freund, Hefter, Homberg, and Reiners, 1984).

- During experiment data collection, contact with the sensor by the human should be eliminated to avoid any uncertainties of human driven motion.

- The testbed must permit the sensor to be rigidly attached to its motion source to prevent any drift in its position or orientation during testing.

- The testbed should be useable with a variety of different displacement transduction technologies. Therefore, the techniques and hardware need to be easily adaptable to accommodate a variety of sensor types while still adhering to a common test method.

Hardware

The testbed hardware components, along with a typical VE spatial transducer, are shown schematically in Figure 1. The testbed components include an IBM AT personal computer with two serial ports, a programmable motion controller and servo amplifier, and a motorized swing arm.

The swing arm apparatus is depicted to scale in Figure 2. It is driven by a permanent magnet DC servo motor (Model ME-5370, EG&G Torque Systems, Watertown MA) bolted to a rigid acrylic plate and wood frame base. A 1 inch diameter pultruded glass-epoxy solid rod, firmly clamped to the top end of the motor shaft, serves as the rotating swing arm. A polycarbonate mounting block which accepts a variety of adapter plates for any of several common spatial sensors (*i.e.*, Ascension, Polhemus, and Logitech devices) is located at the distal end of the swing arm. Nylon fasteners are used both to clamp the halves of the mounting block together and to secure the sensors. The distance, r, from the spatial sensor to the motor shaft center of rotation is adjustable, either by changing the length of glass-epoxy rod, or by sliding the mounting block along the rod. Padded limit stops fixed to the frame restrict swing arm angle, θ, to a maximum range of ±40 degrees. The total weight of the swing arm apparatus is approximately 75 lbs.

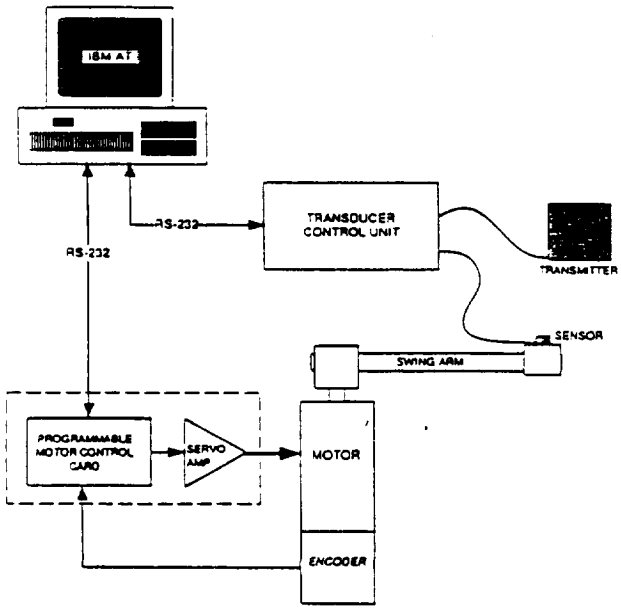

Figure 1. Testbed hardware components.

During tests with electromagnetic devices, the transmitter was secured to the top of a sturdy nonmetallic base (an upside down 35 gallon plastic garbage can). The sensor was attached to the top surface of the mounting block and the ±40 degree swing region faced toward the transmitter. Referring to Figure 2, the height difference between the centers, z_o, ranged from 13.3 to 14.2 inches, dependent on sensor and transmitter size. In all of our experiments with electromagnetic transducers, the separation between the center of the sensor and motor axis was set to r = 14.5 inches. The horizontal distance from the center of the transmitter to the center of rotation of the swing arm was fixed at x_o = 26 inches, ensuring that the sensor always remained within an acceptable operating range (17.5 to 22.5 inches) from the transmitter. With this relative placement of the sensor and transmitter, the sensor did not suffer from any interference that could otherwise have been caused by the motor and the underlying metal floor tile in our lab.

The effect of the motor and the rest of the lab environment on the electromagnetic transducers was checked informally during quasi-static displacements of the swing arm. With the transducer in stream mode, x and y planar data were drawn directly to the screen of the host computer and not erased. The

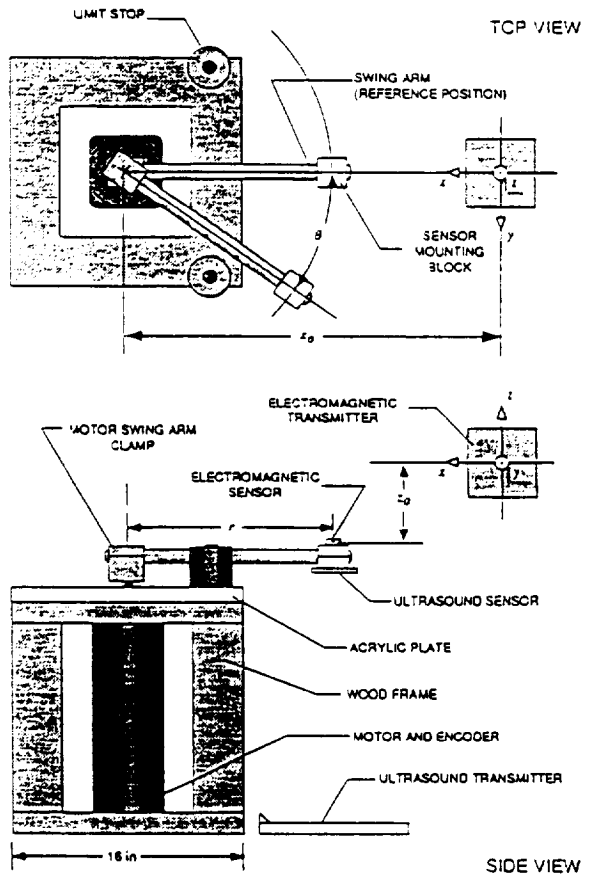

Figure 2. VE spatial sensor testbed. The motorized swing arm apparatus is shown with the electromagnetic transducer coordinate frame. The coordinate frame location is adjusted when the ultrasound based transducer is used.

665 consecutive frames by the available buffer size. The use of the buffer was necessary because writing data directly to disk caused unacceptable time gaps in transducer and encoder motion records.

The spatial transducer always reported its translational displacement, and, depending on the details of the particular test, it would also report its orientation—either as Euler angles, quaternions, direction cosines, or rotation matrix. For the purposes of our subsequent analysis, only the translational output was needed. Thus, to minimize buffer size requirements, the data stored from each frame consisted of only five elements: transducer x, y, and z displacement; transducer time stamp; encoder angle; and encoder time stamp.

EXPERIMENTAL METHOD

The typical experiments for a particular spatial transducer involved testing the device's dynamic response under a variety of its available settings. These transducer settings included report type (*e.g.*, position only, position plus quaternions, position plus Euler angles, *etc.*), and, if available, choices of internal data filter type or the device's work volume. The testing protocol that we followed is outlined here.

At the beginning of the test series with each device, after fastening the sensor to the swing arm and positioning the transmitter with respect to the motor shaft axis, initialization parameters were downloaded to the transducer and the motor control card. Included at this stage were the internal transducer filtering or work volume settings. The transducer's report type along with the spring stiffness to be emulated by the servo system were entered via the keyboard and downloaded at the start of each test in the series. The test series was complete after all combinations of report type and stiffness setting had been examined. Each combination was tested only once for each transducer.

As explained above, the swing arm apparatus was controlled to produce pendulum-like sinusoidal oscillations. The servo's reference position was always set to $\theta = 0$ (corresponding to the transmitter's x axis in Figure 2). At the start of each test, the swing arm was pulled back toward one of the limit stops and then released. Upon the swing arm's release, the data acquisition cycle was initiated from the PC keyboard. Data collection for the test was terminated when the oscillations of the arm had damped out completely.

Data was collected for up to five successive stiffness settings, where each setting was double its predecessor. Since the swing arm is essentially a simple underdamped spring-mass system (*i.e.*, its oscillatory frequency varies in proportion to the square root of spring stiffness), doubling the stiffness increased the oscillation frequency by a factor of 1.4. The frequencies measured at the five increasing settings were approximately 0.95, 1.35, 1.9, 2.7, and 3.8 Hz. Damping in the system, due primarily to motor shaft friction, limited the number of complete oscillations in each test—as few as one cycle at the lowest stiffness and up to 20 or more complete cycles at the highest. Though greater stiffness settings are available in our apparatus, they result in servo amplifier saturations, which in turn trigger unstable (increasing rather than decreasing amplitude) nonlinear (nonconstant frequency) oscillations.

DATA ANALYSIS

In quantifying spatial transducer response to input oscillations, the encoder readings, due to their well documented angular accuracy, resolution, and latency specifications, are treated as a precise measurement of swing arm position.

The analysis of dynamic response in this work is based on the time domain comparison of swing arm angle, ϕ, reconstructed from transducer displacement data against the "true" swing arm angle, θ, measured by the encoder. In principle, since the rotation of the sensor is the same as the swing arm to which it is attached, ϕ, could be read directly from the *orientation* output of the transducer. The intent for these experiments, however, was to examine *displacement* response under a variety of conditions—conditions that could include time lag due to the burden of the transducer's internal computation of orientation format (*i.e.*, Euler angles, quaternions, rotation matrices, or direction cosines).

Data processing required initial conversion of both encoder and spatial sensor readings to a common physical dimension—swing arm angle in degrees. The encoder output, already in angular form, only had to be rescaled. Under the assumption that swing arm motion lay in the x-y plane shown in Figure 2, the transducer's raw x-y outputs were rescaled from transducer units into inches and corrected for offsets between the transmitter and testbed motor shaft. The transducer reconstruction was then calculated from

$$\phi = \tan^{-1}\left(\frac{Y_t}{X_t}\right), \qquad (1)$$

where X_t and Y_t are the rescaled and offset-corrected transducer displacements.

Typical θ and ϕ time histories from experiments with a Polhemus Isotrak sensor (Model 3SI0002, Polhemus Navigation Sciences, Colchester VT) are plotted in Figure 4. The circles denote the discrete time stamped sample points that make up each data record. It can be seen in the plots that the spatial transducer reconstruction always lags the encoder angular measurement. As expected, when internal transducer filtering is increased by switching from "normal" (top plot) to "quiet" mode (bottom plot), the time lag grows and the amplitude of the reconstructed angle, ϕ, decreases relative to θ.

The difference between times for corresponding encoder and spatial transducer at zero crossings was used to estimate the delay in transducer response to displacement inputs. Since individual data samples are unlikely ever to be exactly zero valued, higher resolution estimates of zero crossing times were made feasible by linearly interpolating between sample points on either side of the zero crossing. Linear interpolation is justifiable near the zero crossing because the local slope of sinusoids is nearly constant in this region. The statistical confidence of the lag estimates was improved by measuring and averaging the time differences for both positive and negative zero crossings for all oscillations in a test record. This statistical method is not useable in the lowest servo stiffness cases because averaging cannot be performed with only a single oscillatory cycle per test.

Amplitude fidelity of the transducer in reconstructing swing arm angle was quantified in terms of the average gain—the ratio of average oscillation amplitude—between the two data records, θ and ϕ, for each test. The root mean squared (RMS) magnitude for the entire length of each record was calculated, and the RMS value of the process output, ϕ, was divided by the RMS value of the input, θ, to produce the single gain number for each test.

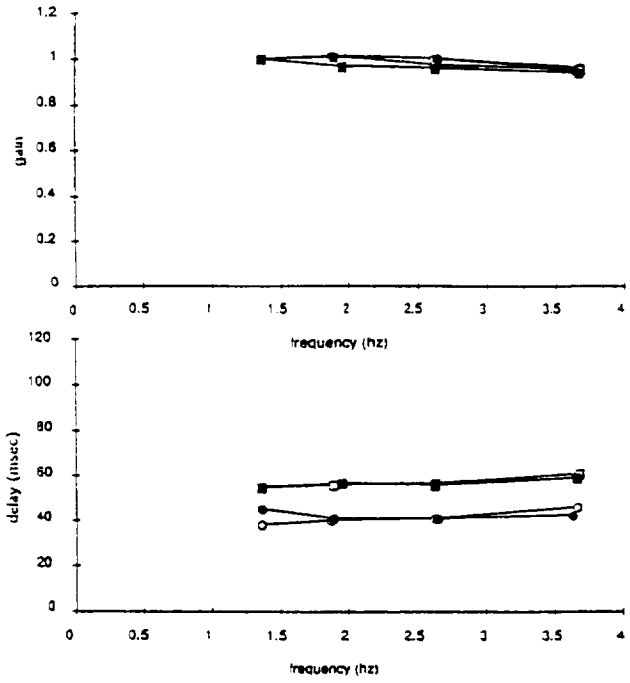

Figure 5. Ascension Bird (serial no. 0059) transducer response. All internal filtering was disabled for data shown by filled markers. Unfilled markers are for data with default low pass and wide AC filter active. Report type: (●,○) position only; (■,□) position plus rotation matrix.

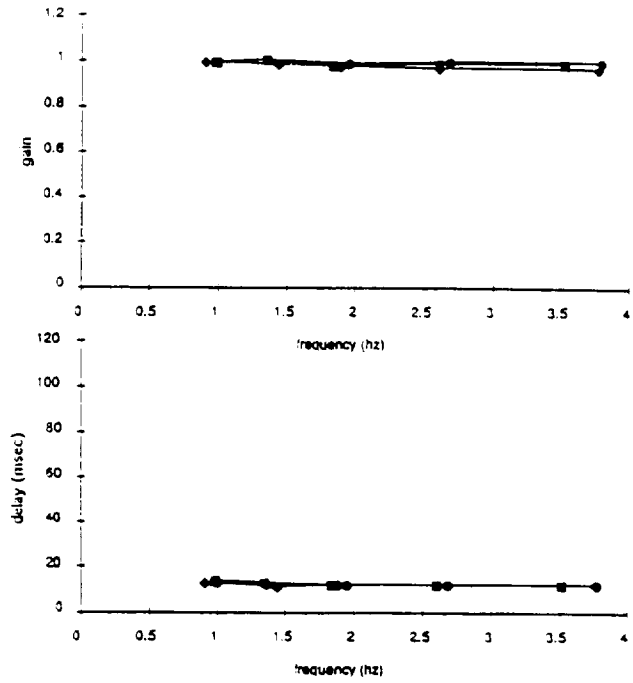

Figure 7. Polhemeus Tracker (serial no. 0026) transducer response. The device has a custom EEPROM that eliminates all internal filtering. Report type: (●) position only; (◆) position plus quaternions: (■) position plus direction cosine matrix.

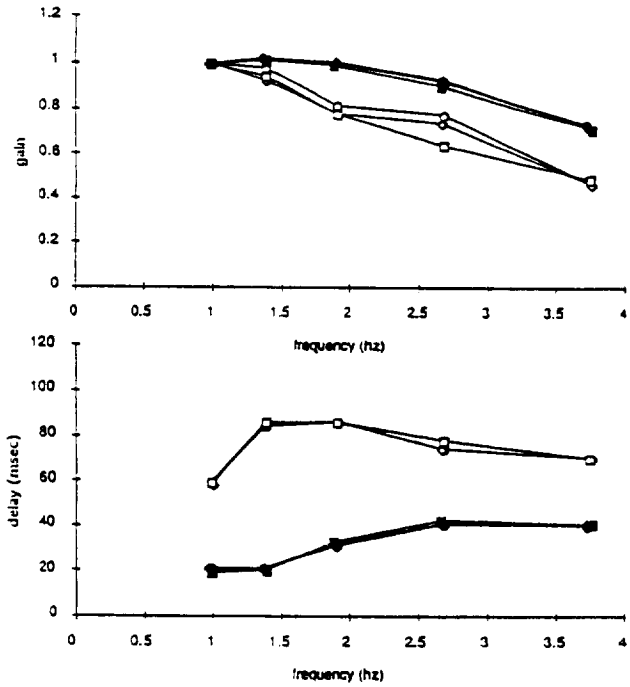

Figure 6. Polhemus Isotrak (serial no. 0436) transducer response. "Normal" mode data are shown with filled markers; "quiet" mode with unfilled markers. Report type: (●,○) position only; (◆,◇) position plus quaternions: (■,□) position plus direction cosine matrix.

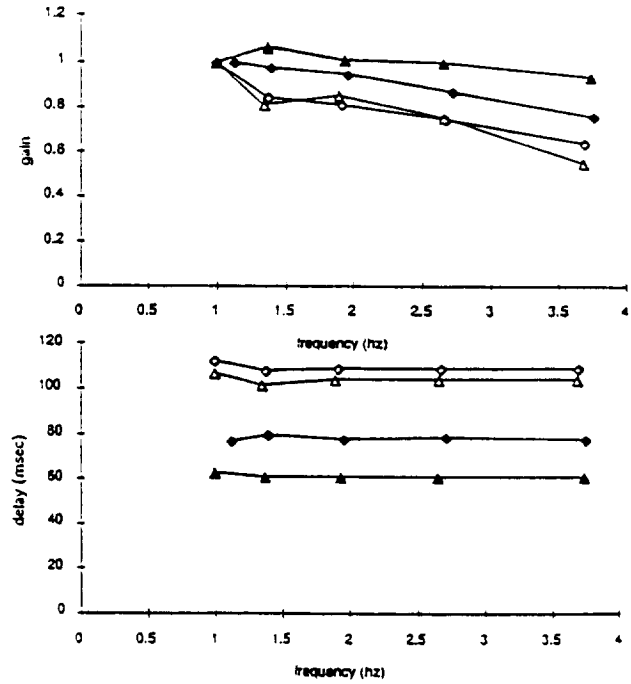

Figure 8. Logitech "Developer's Kit" (serial no. LU00033) transducer response. "Mouse" mode (small work volume) data are shown with filled markers. "Head tracker" mode (large work volume) data are shown with unfilled markers. Report type: (▲,△) position plus Euler angles: (◆,◇) position plus quaternions.